



UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
10/722,614	11/26/2003	Ronald S. Karr	5760-09200	7091

35690 7590 11/21/2006

MEYERTONS, HOOD, KIVLIN, KOWERT & GOETZEL, P.C.
700 LAVACA, SUITE 800
AUSTIN, TX 78701

EXAMINER

TSAI, SHENG JEN

ART UNIT	PAPER NUMBER
----------	--------------

2186

DATE MAILED: 11/21/2006

Please find below and/or attached an Office communication concerning this application or proceeding.

Office Action Summary

Application No.

10/722,614

Applicant(s)

KARR ET AL.

Examiner

Sheng-Jen Tsai

Art Unit

2186

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --

Period for Reply

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE 3 MONTH(S) OR THIRTY (30) DAYS, WHICHEVER IS LONGER, FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133). Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

Status

- 1) ☒ Responsive to communication(s) filed on 16 October 2006.
- 2a) ☒ This action is **FINAL**. 2b) ☐ This action is non-final.
- 3) ☐ Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

Disposition of Claims

- 4) ☒ Claim(s) 1-28 is/are pending in the application.
- 4a) Of the above claim(s) _____ is/are withdrawn from consideration.
- 5) ☐ Claim(s) _____ is/are allowed.
- 6) ☐ Claim(s) _____ is/are rejected.
- 7) ☐ Claim(s) _____ is/are objected to.
- 8) ☐ Claim(s) _____ are subject to restriction and/or election requirement.

Application Papers

- 9) ☐ The specification is objected to by the Examiner.
- 10) ☐ The drawing(s) filed on _____ is/are: a) ☐ accepted or b) ☐ objected to by the Examiner.
Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).
Replacement drawing sheet(s) including the correction is required if the drawing(s) is objected to. See 37 CFR 1.121(d).
- 11) ☐ The oath or declaration is objected to by the Examiner. Note the attached Office Action or form PTO-152.

Priority under 35 U.S.C. § 119

- 12) ☐ Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).
- a) ☐ All b) ☐ Some * c) ☐ None of:
1. ☐ Certified copies of the priority documents have been received.
 2. ☐ Certified copies of the priority documents have been received in Application No. _____.
 3. ☐ Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).

* See the attached detailed Office action for a list of the certified copies not received.

Attachment(s)

- 1) ☐ Notice of References Cited (PTO-892)
- 2) ☐ Notice of Draftsperson's Patent Drawing Review (PTO-948)
- 3) ☒ Information Disclosure Statement(s) (PTO/SB/08)
Paper No(s)/Mail Date _____

- 4) ☐ Interview Summary (PTO-413)
Paper No(s)/Mail Date. _____
- 5) ☐ Notice of Informal Patent Application
- 6) ☐ Other: _____

DETAILED ACTION

1. This Office Action is taken in response to Applicants' Amendment and Remarks filed on July 20, 2006 regarding application 10,722,614 filed on November 26, 2003.

2. Claims 1, 14, 27 and 28 have been amended.

Claims 1-28 are pending for consideration.

3. ***Response to Remarks and Amendments***

Applicants' amendments and remarks have been fully and carefully considered, with the Examiner's response set forth below.

Each of independent claims 1, 14, 27 and 28 has been amended with the additional limitation of "determining a metadata format usable to access data stored on the at least one storage device under a first operating system, wherein the metadata format is determined in response to a request by a host computer system to access the data, and wherein the metadata format is determined based on the host computer system running the first operating system."

Applicants contend that the reference (Rajan et al., US Patent Application Publication 2004/0030822) does not teach this newly added limitation because, in the invention of Rajan et al., the metadata format for the data on the storage appliance is fixed according to the requirements of the storage operating system installed on the storage appliance, and the metadata format is determined when the storage operating system is installed before any request by a host computer system to access the data. The Examiner disagrees with this assessment due to the following reasons:

First, the invention of Rajan et al. is directed toward a storage virtualization selection technique intended for a multi-protocol storage appliance to handle clients (i.e., host computers) running different types of operating systems [abstract]. Figure 1 of Rajan et al. explicitly shows two clients, one running WINDOWS operating system [figure 1, 160a] and the other running UNIX operating system [figure 1, 160b].

Second, since the host computers use different operating systems with different storage formats, these differences are reflected in the corresponding operating system metadata, as explicitly disclosed by Rajan et al. [For example, a client 160a running the Windows operating system may communicate with the storage appliance 100 using the Common Internet File System (CIFS) protocol over TCP/IP. On the other hand, a client 160b running the UNIX operating system may communicate with the multi-protocol appliance using either the Network File System (NFS) protocol over TCP/IP or the Direct Access File System (DAFS) protocol over a virtual interface (VI) transport in accordance with a remote DMA (RDMA) protocol over TCP/IP. It will be apparent to those skilled in the art that other clients running other types of operating systems may also communicate with the integrated multi-protocol storage appliance using other file access protocols (paragraph 0026)].

Third, figure 2 of Rajan et al. shows, as the Applicants observe, that both CIFS [for WINDOW operating system, see above] and NFS [for UNIX operating system, see above] are installed and made available during the installation phase. However, the determination as to which one of the CIFS and NFS should be used is made dynamically depending on which client [figure 1, 160a or 160b] initiates the access

Art Unit: 2186

request [The storage adapter 128 cooperates with the storage operating system 200 executing on the storage appliance to access information requested by the clients (paragraph 0030)]. In other words, although the installation of both CIFS and NFS is fixed, the determination of which one to use is made dynamically.

Therefore, the Examiner's position regarding the status of claims 1, 14, 19 and 24, and those claims dependent from them, remain the same as stated in the previous Office Action.

Another iteration of claim analysis has been made in response to the amendments. Refer to the corresponding sections of the following claim analysis for details.

Claim Rejections - 35 USC § 102

(e) the invention was described in (1) an application for patent, published under section 122(b), by another filed in the United States before the invention by the applicant for patent or (2) a patent granted on an application for patent by another filed in the United States before the invention by the applicant for patent, except that an international application filed under the treaty defined in section 351(a) shall have the effects for purposes of this subsection of an application filed in the United States only if the international application designated the United States and was published under Article 21(2) of such treaty in the English language.

4. Claims 1-28 are rejected under 35 U.S.C. 102(e) as being anticipated by Rajan et al. (U.S. Patent Application Publication 2004/0030822).

As to claim 1, Rajan et al. disclose a **storage subsystem** [Storage Virtualization by Layering Virtual Disk Objects on a File System (title); figure 1], **comprising:**
at least one storage device [figure 1, 150 shows a plurality of storage disks]; **and**
a storage virtualization controller [the corresponding storage virtualization controller is the multi-protocol storage appliance unit shown in figure 1, 100, comprising a processor (122), a memory (124) with storage operating system (200), a network

Art Unit: 2186

adapter (125), a storage adapter (128) and a network target adapter (126); figure 2],
wherein the storage virtualization controller is communicatively coupled to the at least one storage device [as shown in figure 1 via the storage adapter (128)], and wherein the storage virtualization controller is operable to:

determining a metadata format usable to access data stored on the at least one storage device under a first operating system, wherein the metadata format is determined in response to a request by a host computer system to access the data, and wherein the metadata format is determined based on the host computer system running the first operating system [First, the invention of Rajan et al. is directed toward a storage virtualization selection technique intended for a multi-protocol storage appliance to handle clients (i.e., host computers) running different types of operating systems (abstract). Figure 1 of Rajan et al. explicitly shows two clients, one running WINDOWS operating system (figure 1, 160a) and the other running UNIX operating system (figure 1, 160b); Second, since the host computers use different operating systems with different storage formats, these differences are reflected in the corresponding operating system metadata, as explicitly disclosed by Rajan et al. [For example, a client 160a running the Windows operating system may communicate with the storage appliance 100 using the Common Internet File System (CIFS) protocol over TCP/IP. On the other hand, a client 160b running the UNIX operating system may communicate with the multi-protocol appliance using either the Network File System (NFS) protocol over TCP/IP or the Direct Access File System (DAFS) protocol over a virtual interface (VI) transport in accordance with a remote DMA (RDMA) protocol over

TCP/IP. It will be apparent to those skilled in the art that other clients running other types of operating systems may also communicate with the integrated multi-protocol storage appliance using other file access protocols (paragraph 0026)]; Third, figure 2 of Rajan et al. shows, as the Applicants observe, that both CIFS (for WINDOW operating system, see above) and NFS (for UNIX operating system, see above) are installed and made available during the installation phase. However, the determination as to which one of the CIFS and NFS should be used is made dynamically depending on which client (figure 1, 160a or 160b) initiates the access request (The storage adapter 128 cooperates with the storage operating system 200 executing on the storage appliance to access information requested by the clients (paragraph 0030)). In other words, although the installation of both CIFS and NFS is fixed, the determination of which one to use is made dynamically; figures 4 and 5 show examples of the metadata].

generate operating system metadata in accordance with the determined metadata format [The storage adapter 128 cooperates with the storage operating system 200 executing on the storage appliance to access information requested by the clients (paragraph 0030); For example, a client 160a running the Windows operating system may communicate with the storage appliance 100 using the Common Internet File System (CIFS) protocol over TCP/IP. On the other hand, a client 160b running the UNIX operating system may communicate with the multi-protocol appliance using either the Network File System (NFS) protocol over TCP/IP or the Direct Access File System (DAFS) protocol over a virtual interface (VI) transport in accordance with a remote DMA (RDMA) protocol over TCP/IP. It will be apparent to those skilled in the

Art Unit: 2186

art that other clients running other types of operating systems may also communicate with the integrated multi-protocol storage appliance using other file access protocols (paragraph 0026); since the host computers use different operating systems with different storage formats, these differences are reflected in the corresponding operating system metadata; figure 4, 410 shows the metadata generated, including type (412), size (414), time stamps (416), UID (418), GID (420) and Xinode (430); figure 5, 501, 512, 522, 542, 552 and 562 provide more examples of metadata] **for the at least one storage device** [figure 1, 150 shows a plurality of storage disks], **wherein the operating system metadata emulates a storage volume hosted under the first operating system** [figure 3 shows the storage volume information indicated by VDISK module (330) including LUN (Logical Unit Number), IGROUP and Map Binding; as used herein, the term "storage operating system" generally refers to the computer-executable code operable on a computer that manages data access and may, in the case of a multi-protocol storage appliance, implement data access semantics, such as the Data ONTAP storage operating system, which is implemented as a microkernel. The storage operating system can also be implemented as an application program operating over a general-purpose operating system, such as UNIX or Windows NT, or as a general-purpose operating system with configurable functionality, which is configured for storage applications as described herein (paragraph 0035)]; **and send the operating system metadata to the host computer system** [the corresponding host computer system is the client systems shown in figure 1, 160a and 160b], **wherein the operating system metadata enables the host computer system**

to recognize the storage device as the storage volume hosted under the first operating system [Whereas clients of a NAS-based network environment have a storage viewpoint of files, the clients of a SAN-based network environment have a storage viewpoint of blocks or disks. To that end, the multi-protocol storage appliance 100 presents (exports) disks to SAN clients through the creation of logical unit numbers (luns) or vdisk objects. A vdisk object (hereinafter "vdisk") is a special file type that is implemented by the virtualization system and translated into an emulated disk as viewed by the SAN clients. The multi-protocol storage appliance thereafter makes these emulated disks accessible to the SAN clients through controlled exports, as described further herein (paragraph 0023)].

As to claim 2, Rajan et al. teach that **the operating system metadata enables a block storage I/O stack in the first operating system on the host computer system to recognize the storage device as a partition** [To that end, the multi-protocol storage appliance 100 presents (exports) disks to SAN clients through the creation of logical unit numbers (luns) or vdisk objects. A vdisk object (hereinafter "vdisk") is a special file type that is implemented by the virtualization system and translated into an emulated disk as viewed by the SAN clients. The multi-protocol storage appliance thereafter makes these emulated disks accessible to the SAN clients through controlled exports, as described further herein (paragraph 0023). Note that logical unit numbers (luns) and vdisk objects are both special forms of partition; figure 3 shows that data being partitioned into the form of "record" (372)].

As to claim 3, Rajan et al. teach that **the operating system metadata enables a block storage I/O stack in the first operating system on the host computer system to recognize the storage device as a host-virtual object** [To that end, the multi-protocol storage appliance 100 presents (exports) disks to SAN clients through the creation of logical unit numbers (luns) or vdisk objects. A vdisk object (hereinafter "vdisk") is a special file type that is implemented by the virtualization system and translated into an emulated disk as viewed by the SAN clients. The multi-protocol storage appliance thereafter makes these emulated disks accessible to the SAN clients through controlled exports, as described further herein (paragraph 0023). Note that data is viewed by the clients as virtual disk (vdisk) object].

As to claim 4, Rajan et al. teach that **the operating system metadata enables a driver on the host computer system to recognize the storage device as an enclosed volume** [To that end, the multi-protocol storage appliance 100 presents (exports) disks to SAN clients through the creation of logical unit numbers (luns) or vdisk objects], **wherein the driver is layered above a block storage I/O stack in the first operating system** [The storage operating system comprises a series of software layers organized to form an integrated network protocol stack ... (paragraph 0037)].

As to claim 5, Rajan et al. teach that **the storage virtualization controller is operable to configure the operating system metadata in response to a requirement of the first operating system** [The vdisk is thereafter created as a storage object within a volume and, thus, inherits the underlying reliability configuration associated with that volume (abstract); The file server, or filer, may be further

configured to operate according to a client/server model of information delivery to thereby allow many client systems (clients) to access shared resources, such as files, stored on the filer (paragraph 0003)].

As to claim 6, Rajan et al. teach that a **management environment is configured to supply operating system types and operating system metadata configuration requirements to the storage virtualization controller** [The file server, or filer, may be further configured to operate according to a client/server model of information delivery to thereby allow many client systems (clients) to access shared resources, such as files, stored on the filer (paragraph 0003)], **wherein the operating system types comprise the first operating system** [The storage operating system can also be implemented as an application program operating over a general-purpose operating system, such as UNIX or Windows NT, or as a general-purpose operating system with configurable functionality, which is configured for storage applications as described herein (paragraph 0035); figure 1 shows that client 160a runs WINDOWS operating system and client 160b runs UNIX operating system].

As to claim 7, Rajan et al. teach that **in generating the operating system metadata for the storage device, the storage virtualization controller is operable to add a storage property to identify an offset and a length of the storage volume** [figure 4, 410 shows the metadata generated, including type (412), size (414), time stamps (416), UID (418), GID (420) and Xinode (430); figure 3 shows the storage volume information indicated by VDISK module (330) including LUN (Logical Unit Number), IGROUP and Map Binding].

As to claim 8, Rajan et al. teach that **an operation is provided to configure operating system types and operating system metadata configuration requirements for generating the operating system metadata, wherein the operating system types comprise the first operating system** [The vdisk is thereafter created as a storage object within a volume and, thus, inherits the underlying reliability configuration associated with that volume (abstract); The file server, or filer, may be further configured to operate according to a client/server model of information delivery to thereby allow many client systems (clients) to access shared resources, such as files, stored on the filer (paragraph 0003)].

As to claim 9, Rajan et al. teach that **the storage virtualization controller is operable to receive user input to select one of a plurality of operating system types for the operating system metadata, wherein the operating system types comprise the first operating system** [The file server, or filer, may be further configured to operate according to a client/server model of information delivery to thereby allow many client systems (clients) to access shared resources, such as files, stored on the filer (paragraph 0003)].

As to claim 10, Rajan et al. teach that **the storage virtualization controller is operable to send an operating system metadata configuration instruction to the storage device through a vendor-unique I/O request to the storage device** [Whereas clients of a NAS-based network environment have a storage viewpoint of files, the clients of a SAN-based network environment have a storage viewpoint of blocks or disks ; To that end, the multi-protocol storage appliance 100 presents

(exports) disks to SAN clients through the creation of logical unit numbers (luns) or vdisk objects. A vdisk object (hereinafter "vdisk") is a special file type that is implemented by the virtualization system and translated into an emulated disk as viewed by the SAN clients. The multi-protocol storage appliance thereafter makes these emulated disks accessible to the SAN clients through controlled exports, as described further herein (paragraph 0023); figure 4, 410 shows the metadata generated, including type (412), size (414), time stamps (416), UID (418), GID (420) and Xinode (430)].

As to claim 11, Rajan et al. teach that **the operating system metadata emulates a storage volume hosted under a first operating system and one or more additional operating systems** [the term "storage operating system" generally refers to the computer-executable code operable on a computer that manages data access and may, in the case of a multi-protocol storage appliance, implement data access semantics, such as the Data ONTAP storage operating system, which is implemented as a microkernel. The storage operating system can also be implemented as an application program operating over a general-purpose operating system, such as UNIX or Windows NT, or as a general-purpose operating system with configurable functionality, which is configured for storage applications as described herein (paragraph 0035); figure 1 shows that client 160a runs WINDOWS operating system and client 160b runs UNIX operating system]; **and wherein the operating system metadata enables a layered driver on the host computer system to recognize the storage device** [The storage operating system comprises a series of

software layers organized to form an integrated network protocol stack ... (paragraph 0037)].

As to claim 12, Rajan et al. teach **using a layered driver on the host computer system to provide access to a storage volume mapped within a Logical Unit, wherein the Logical Unit is provided by an external device or an external virtualization layer** [figure 3 shows the storage volume information indicated by VDISK module (330) including LUN (Logical Unit Number), IGROUP and Map Binding; To that end, the multi-protocol storage appliance 100 presents (exports) disks to SAN clients through the creation of logical unit numbers (luns) or vdisk objects (paragraph 0023); The storage operating system comprises a series of software layers organized to form an integrated network protocol stack ... (paragraph 0037)].

As to claim 13, Rajan et al. teach that **a management environment is configured to supply a preferred name of the storage device to software on the host computer system** [a storage operating system 200 that provides a virtualization system (and, in particular, a file system) to logically organize the information as a hierarchical structure of named directory, file and virtual disk (vdisk) storage objects on the disks 130 (paragraph 0022)].

As to claim 14, refer to "As to claim 1" presented earlier in this Office Action.

As to claim 15, refer to "As to claim 2" presented earlier in this Office Action.

As to claim 16, refer to "As to claim 3" presented earlier in this Office Action.

As to claim 17, refer to "As to claim 4" presented earlier in this Office Action.

As to claim 18, refer to "As to claim 5" presented earlier in this Office Action.

As to claim 19, refer to "As to claim 6" presented earlier in this Office Action.

As to claim 20, refer to "As to claim 7" presented earlier in this Office Action.

As to claim 21, refer to "As to claim 8" presented earlier in this Office Action.

As to claim 22, refer to "As to claim 9" presented earlier in this Office Action.

As to claim 23, refer to "As to claim 10" presented earlier in this Office Action.

As to claim 24, refer to "As to claim 11" presented earlier in this Office Action.

As to claim 25, refer to "As to claim 12" presented earlier in this Office Action.

As to claim 26, refer to "As to claim 13" presented earlier in this Office Action.

As to claim 27, refer to "As to claim 1" presented earlier in this Office Action.

Further, figure 6 of Rajan et al. shows the flowchart of the computer programs that implement the storage operating system.

As to claim 28, refer to "As to claim 1" presented earlier in this Office Action.

5. Related Prior Art of Record

The following list of prior art is considered to be pertinent to applicant's invention, but not relied upon for claim analysis conducted above.

- Markson et al., (US Patent Application Publication 2002/0103889), "Virtual Storage Layer Approach for Dynamically Associating Computer Storage with Processing Hosts."
- Oliveira et al., (US 6,889,309), "Method and Apparatus for Implementing an Enterprise Virtual Storage System."
- Idei et al., (US Patent Application Publication 2003/0177330), "Computer System."
- Takamoto et al., (US 6,792,557), "Storage Area Network system."
- Glider, (US 7,020,760), "Hybrid Logical Block Virtualization System for a Storage Area Network."

Conclusion

6. Claims 1-28 are rejected as explained above.

7. **THIS ACTION IS MADE FINAL.** Applicant is reminded of the extension of time policy as set forth in 37 CFR 1.136(a).

A shortened statutory period for reply to this final action is set to expire THREE MONTHS from the mailing date of this action. In the event a first reply is filed within TWO MONTHS of the mailing date of this final action and the advisory action is not mailed until after the end of the THREE-MONTH shortened statutory period, then the shortened statutory period will expire on the date the advisory action is mailed, and any extension fee pursuant to 37 CFR 1.136(a) will be calculated from the mailing date of the advisory action. In no event, however, will the statutory period for reply expire later than SIX MONTHS from the mailing date of this final action.

Art Unit: 2186


8. Any inquiry concerning this communication or earlier communications from the examiner should be directed to Sheng-Jen Tsai whose telephone number is 571-272-4244. The examiner can normally be reached on 8:30 - 5:00.

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Matthew Kim can be reached on 571-272-4182. The fax phone number for the organization where this application or proceeding is assigned is 571-273-8300.

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see <http://pair-direct.uspto.gov>. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free).

Sheng-Jen Tsai
Examiner
Art Unit 2186

November 11, 2006


PIERRE BATAILLE
PRIMARY EXAMINER
11/16/06